Hochschule für Technik
und Wirtschaft Berlin

University of Applied Sciences

Challenges and concepts for dynamic metadata

Hermann Heßling

PUNCH4NFDI

TA5

March 17, 2022

PUNCH4NFDI - Proposal

4.2 Metadata standards

The data landscape of PUNCH is very diverse, and even more so are the systems used to standardise and record metadata. Many advanced community- or experiment-specific approaches exist, albeit with only limited scope. The NFDI process offers a chance to take a fresh look at issues of **metadata and their use** within the PUNCH community. It quickly becomes apparent that there is **no common understanding of the term "metadata"**. For a productive discussion of metadata standards, the following **distinction** shall be used:

- In-file metadata
  - Astronomy: FITS
  - Particle physics: ROOT
- Data Lake
  - Metadata stores
- …

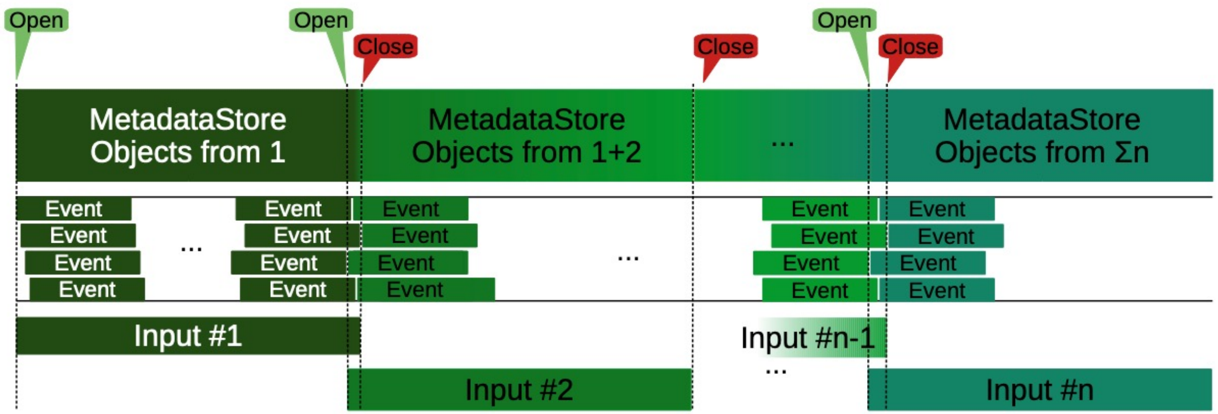

[ATLAS (2021): multi-threaded metadata service]

- General strategy: "extensions or adaptions of metadata schemes"
- Challenge: out-of-file metadata

PUNCH4NDFI - Proposal

5 Work Programme

This section lays out the detailed work programme of the PUNCH4NFDI project

- TA2 – **WP 2.1**: Standardized access to data and **metadata**
- TA4 – **WP 4.2**: Mapping and collating **metadata** („**standard metadata**“)
  - Incompatible metadata schemes: EUDAT, EOSC, Rucio, VO, …
  - Approaches for unifying access to data and metadata
    - Particle physics: CERN open data project
    - Astronomy: IOVA
  - Goal: definition of layers of metadata
    - Top level: publication level
    - Lowest level: raw data [note: in the long-term, (almost) no raw data]
- TA6 – **WP 6.3**: Cross-community efforts towards **FAIR data**
  - **Extended metadata** -> needed for accessing cross-community data
  - Dynamic Metadata -> needed for coping with demands from Dynamic Life Cycle

Dynamic Life Cycle

Data Irreversibility: strategies

- Dynamic filtering
  - Extracting relevant information out. of data streams in realtime
  - Machine learning => smart sensors
- Dynamic archiving
  - Feedback from offline workflows to sensor control in near-realtime
  - Offline ⇒ online computing
- Scaling
  - Online (e. g. parallelisation of workflows – proprietary, in general)
  - Offline (e. g. analysis of Data Monster)
- Reproducibility
  - Reconstructing how decisions were taken
  - Simulations are essential for validation and understanding

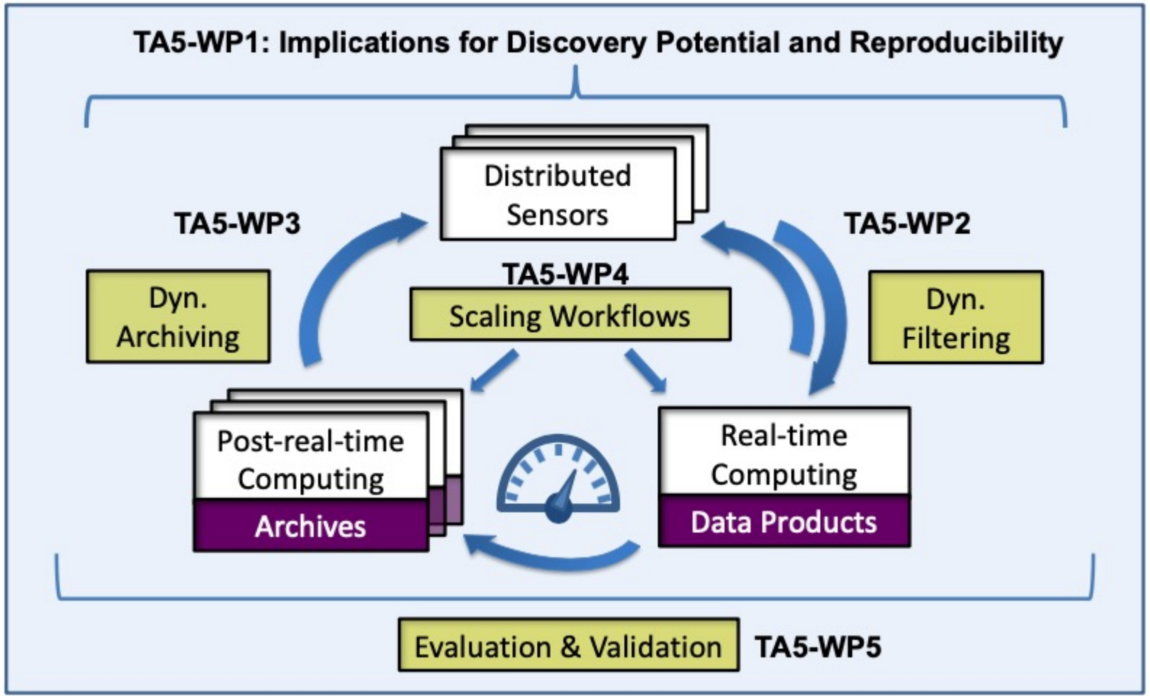

Dynamic Life Cycle Model

Dynamic Life Cycle

Data Irreversibility: strategies

- Dynamic filtering
  - Extracting relevant information out of data streams in realtime
  - Machine learning
- Dynamic archiving
  - Feedback from offline workflows to sensor control in near-realtime
  - Offline ⇒ online computing
- Scaling
  - Online (e. g. parallelisation of workflows – proprietary, in general)
  - Offline (e. g. analysis of Data Monster)
- Reproducibility
  - Reconstructing how decisions were taken
  - Simulations are essential for validation and understanding

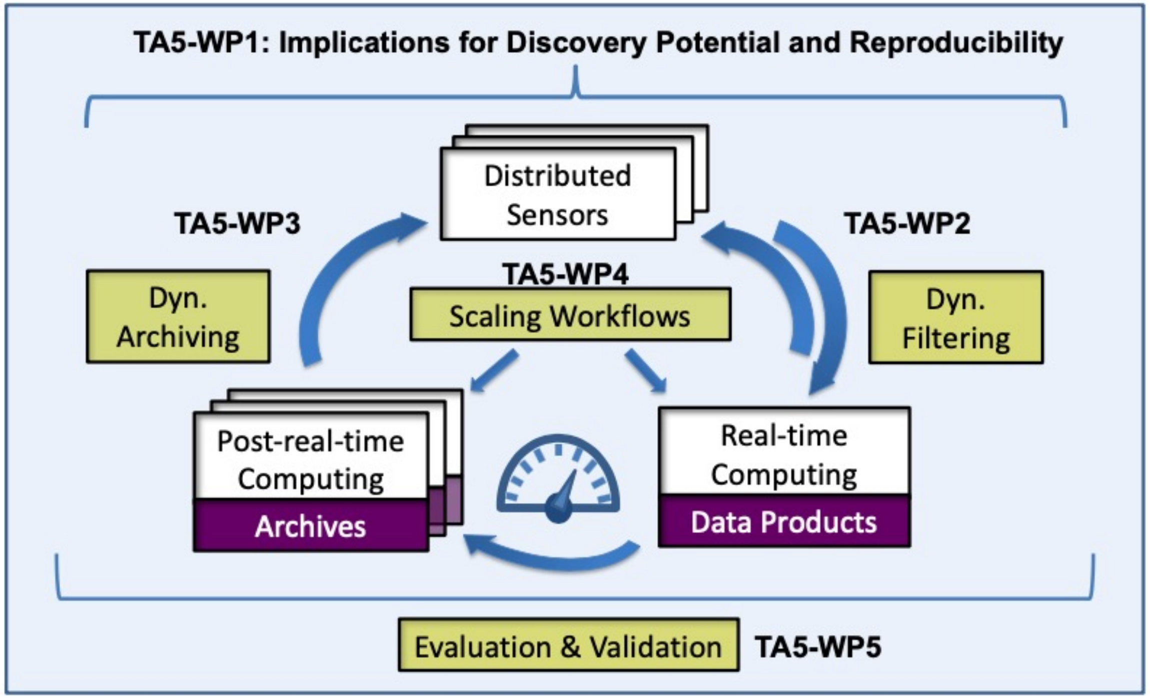

Dynamic Life Cycle Model

Dynamic Metadata

- Drastic increase of metadata volumes
  - Software of Dyn. Filtering (⇒ reproducibility)
  - Constant updates of "quality measures" of archived data
- Flexible data models

Relational vs. NoSQL databases

- Relational database
  - Structured data (tables, normalized data)
  - Query language (SQL)
- Databases with no SLQ
  - Unstructured data
  - Models for accessing and managing data:

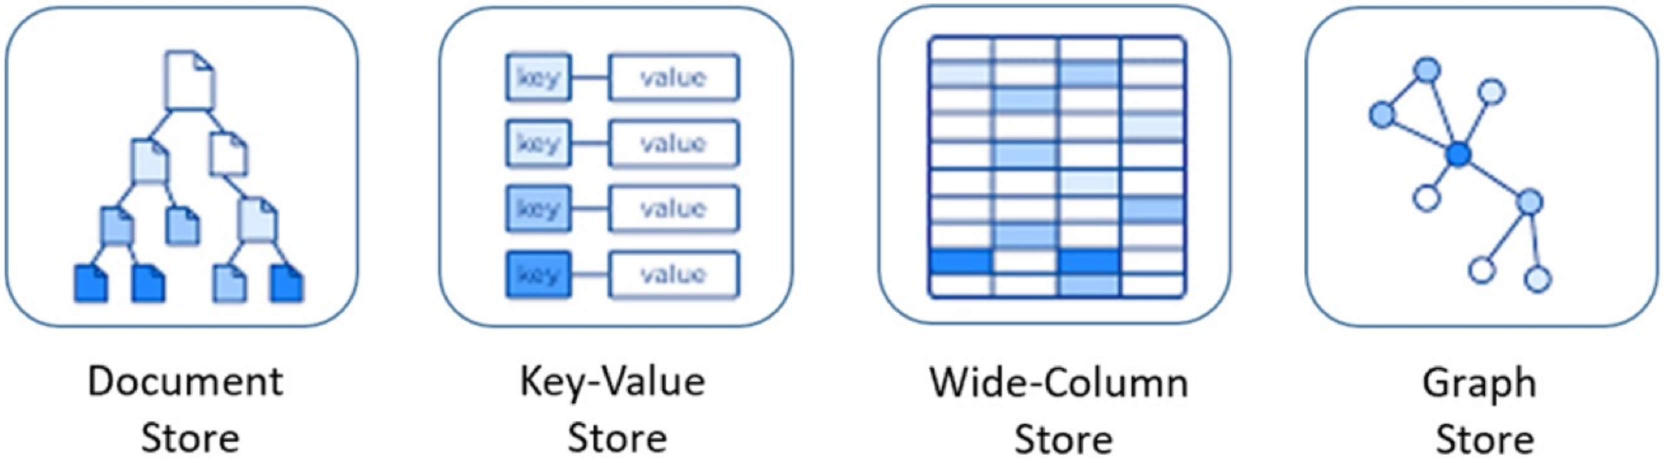

Relational vs. NoSQL databases

- **CAP theorem**
  - A distributed data systems can guarantee only two out of three properties
- **SQL database**
  - Standard: running on one server
    - Copies of DB on secondary servers
  - Good for complex queries
  - **Not designed for change**
    - Modelling in advance, takes up to years
  - **Limited in volume** (see below)
- **NoSQL databases**
  - **Running on (geographically) distributed data** (-> cloud)
  - Good if queries are simple
  - **Huge data volumes**

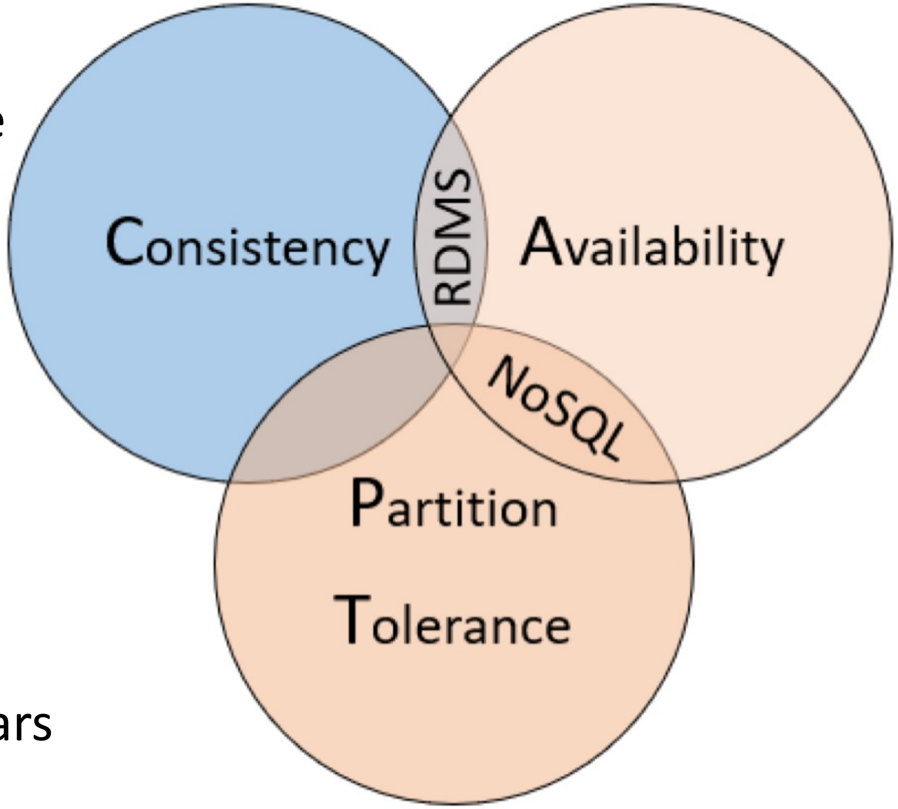

Limits of SQL DBs

	Max DB size	Max table size	Max row size	Max blob size
MySQL	∞	256 TB	64 KB	4 GB
Oracle	2 PB	4 GB	8 K	128 TB
PostgreSQL	**∞**	**32 TB**	**1.6 TB**	**4 TB**
SAP Hana	?	?	?	?
SQLite	128 TB	file size	file size	2 GB

[https://en.wikipedia.org/wiki/Comparison_of_relational_database_management_systems]

Conclusion: SQL DBs do not meet metadata requirements of the Dynamic Life Cycle

Metadata : generic frameworks

- UNICORE
  - HPC middleware service (open source)
  - Accessing distributed computing and federated data resources
- UNICORE metadata system
  - Metadata stored in files close to the data
    - Key-value pairs
  - Extraction: retrieve metadata from files (using Apache Tika)
  - Search: querying metadata (using Apache Lucene)
  - Searching in one store or multiple stores across UNICORE federation

  - Adaption to a particular community: hard work

Metadata : generic frameworks

- Rucio 1.2
  - ATLAS@LHC: highly scalable data management framework
  - Accessing distributed computing and federated data resources

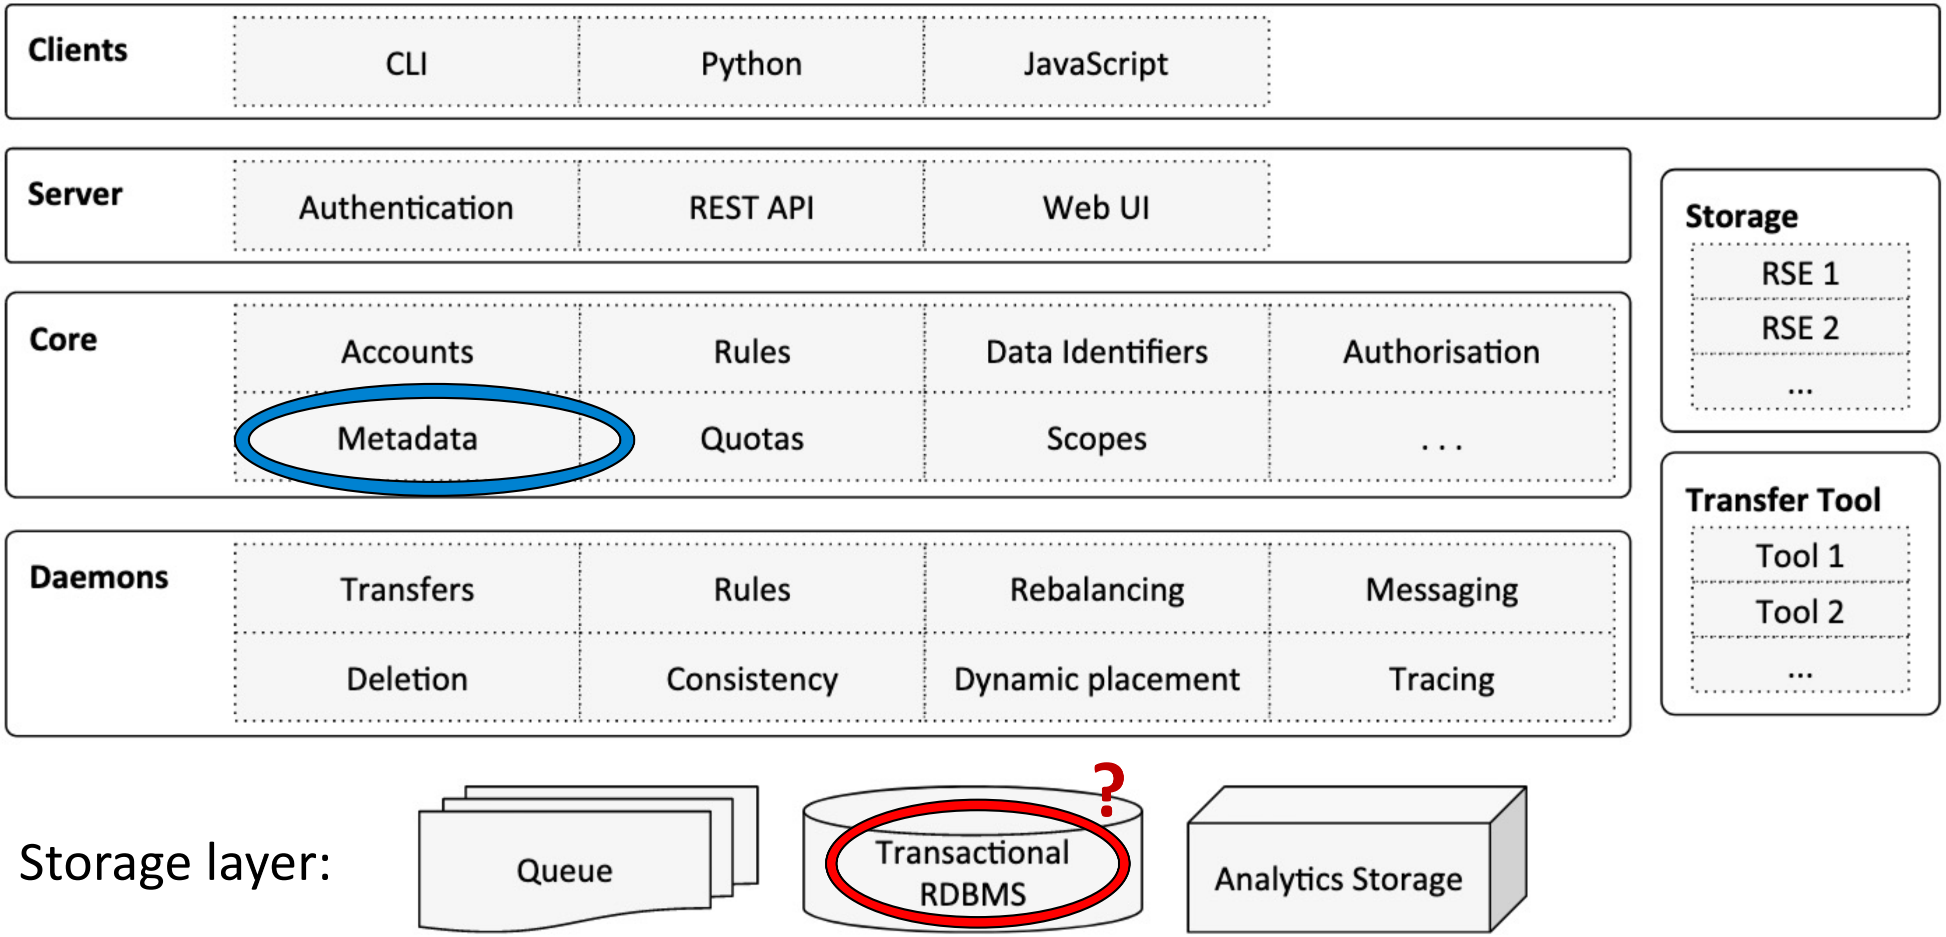

- Future (⇒ astronomy): supporting arbitrary and **flexible metadata**

Metadata : specific frameworks

- XTENS 2
  - Repository for heterogeneous data in life science (biomedicine)
    - Manage heterogenous data (samples, any kind of data)
    - Roles to handle data access for any authenticated user
    - ? Complex queries
  - **Database: PostgreSQL**
    - **Flexible JSON data model**
    - XTENS 1: MySQL, Excel
  - Distributed data: iRODS

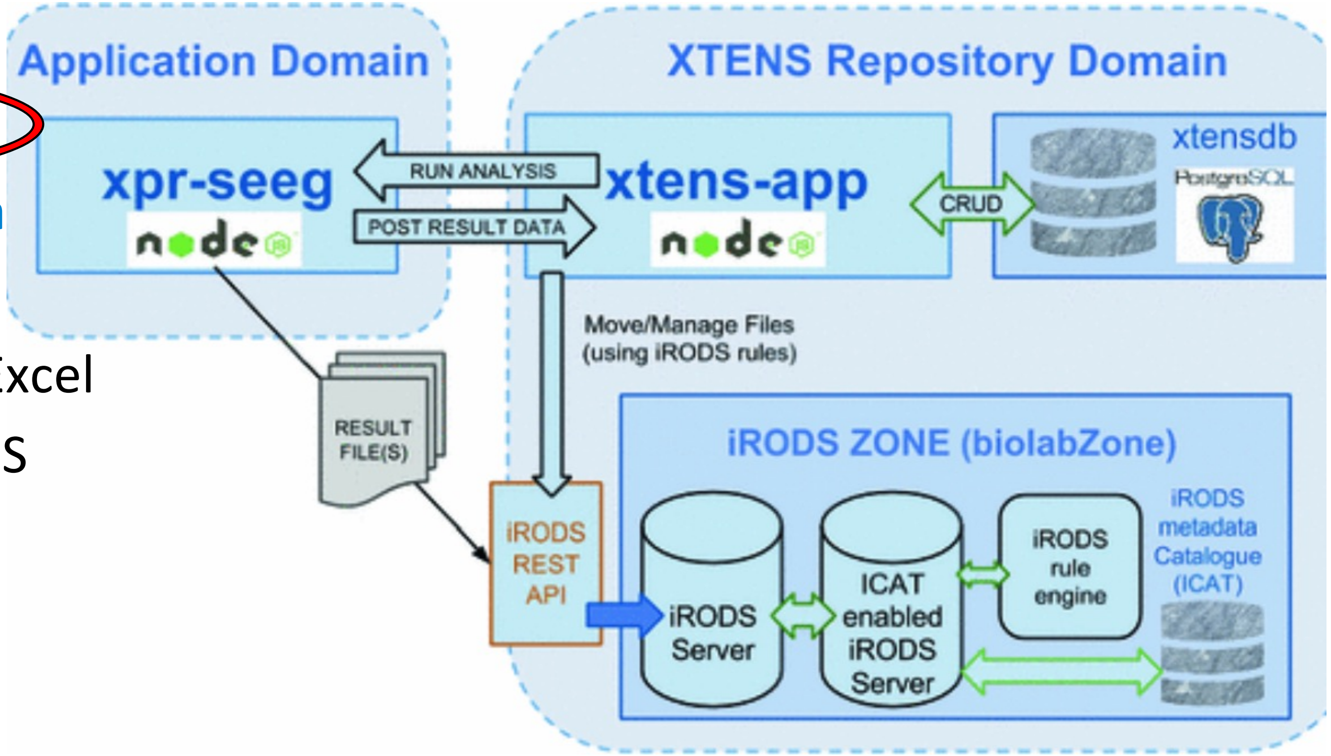

[https://github.com/xtens-suite/xtens-app]

Metadata : specific frameworks

- MASi (Metadata Management for Applied Sciences)
  - Generic metadata programming interface
  - Open Archive Initiative (OAI) protocol for metadata harvesting
  - **Supports multi-community research:** geography, chemistry, digital humanities
  - Multi-model **NoSQL DB**
    - Key-values
    - Documents
    - Graphs
  - **Future (2019):** ? integration into UNICORE

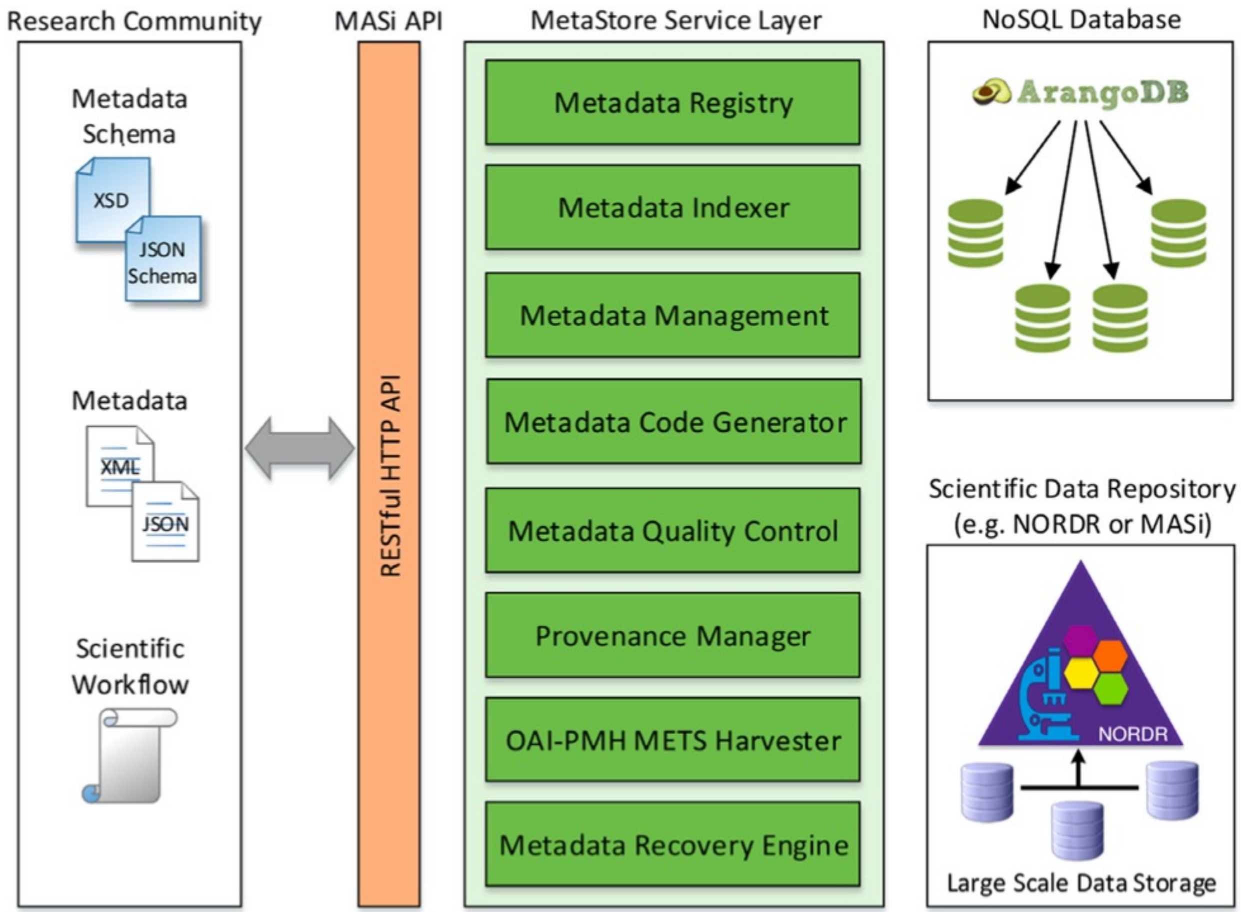

https://**tu-dresden.de/zih**/dienste/service-katalog/zusammenarbeiten-und-forschen/forschungsdatenmanagement/**masi**

Metadata : hybrid systems

ATLAS @ LHC

- Hadoop (NoSQL): full event infos
- Oracle (SQL): reduced infos for faster queries

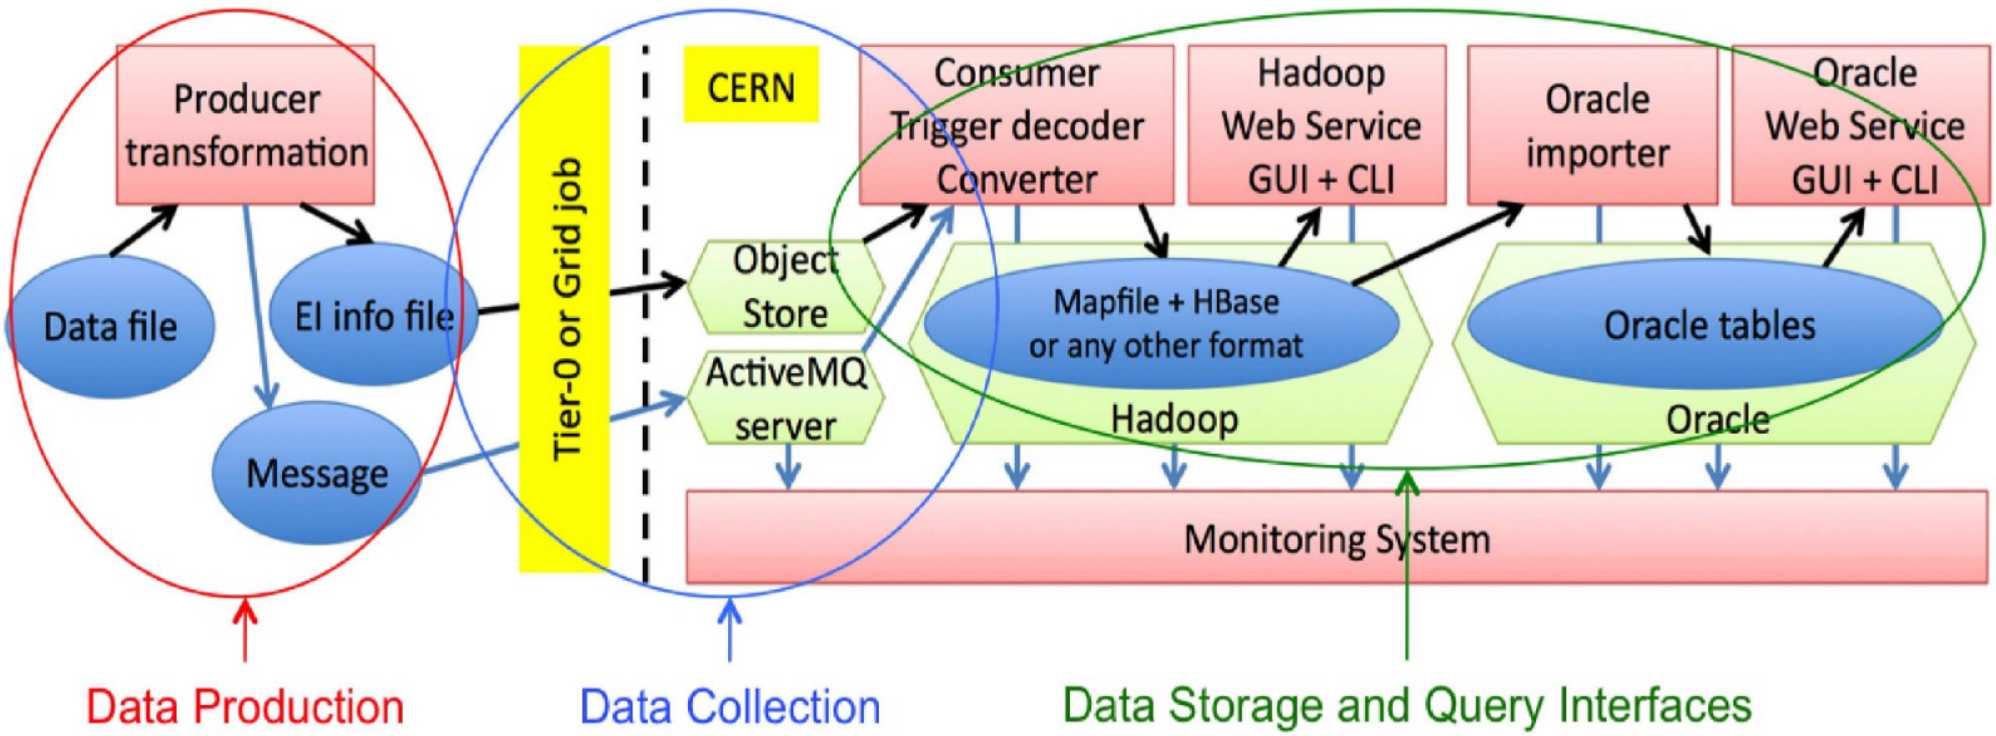

Figure 3: Event Index architecture and data flow.

[2017: http://lamp.ictp.it/index.php/aphysrev/article/download/1598/576]

Metadata : interfaces

ATLAS @ LHC

- AMI (ATLAS Metadata Interface) 2.0
- Metadata aggregation + transformation + cataloging

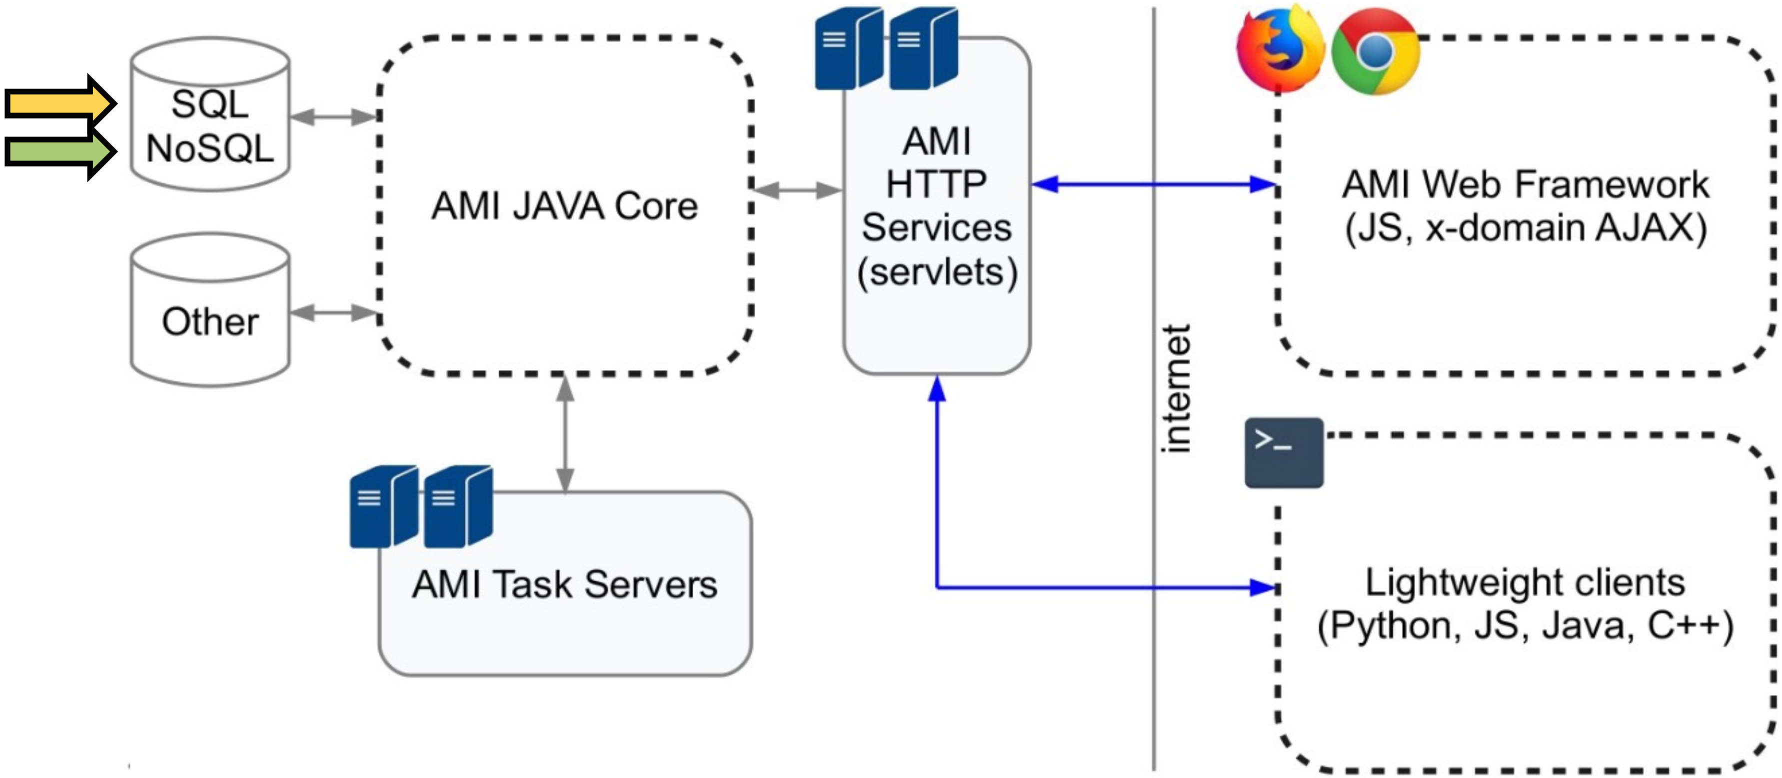

[https://www.epj-conferences.org/articles/epjconf/pdf/2019/19/epjconf_chep2018_05046.pdf]

TA5 ⟺ TA2/3 ⟺ TA4

Thank you !